arXiv:1206.4912v1 [cs.DS] 21 Jun 2012

Preprocessing Subgraph and Minor Problems: When Does a Small Vertex Cover Help?*

Fedor V. Fomin[1], Bart M. P. Jansen[2], and Michał Pilipczuk[1]

[1] Department of Informatics, University of Bergen, N-5020 Bergen, Norway
`{fomin,michal.pilipczuk}@ii.uib.no`
[2] Utrecht University, P.O. Box 80.089, 3508 TB Utrecht, The Netherlands
`b.m.p.jansen@uu.nl`

Abstract. We prove a number of results around kernelization of problems parameterized by a vertex cover of the input graph. We provide two simple general conditions characterizing problems admitting kernels of polynomial size. Our characterizations not only give generic explanations for the existence of many known polynomial kernels for problems like Odd Cycle Transversal, Chordal Deletion, η-Transversal, Long Path, Long Cycle, or H-packing, parameterized by the size of a vertex cover, they also imply new polynomial kernels for problems like $\mathcal{F}$-Minor-Free Deletion, which is to delete at most k vertices to obtain a graph with no minor from a fixed finite set $\mathcal{F}$.

While our characterization captures many interesting problems, the kernelization complexity landscape of problems parameterized by vertex cover is much more involved. We demonstrate this by several results about induced subgraph and minor containment, which we find surprising. While it was known that testing for an induced complete subgraph has no polynomial kernel unless NP $\subseteq$ coNP/poly, we show that the problem of testing if a graph contains a given complete graph on t vertices as a minor admits a polynomial kernel. On the other hand, it was known that testing for a path on t vertices as a minor admits a polynomial kernel, but we show that testing for containment of an induced path on t vertices is unlikely to admit a polynomial kernel.

1 Introduction

Kernelization is an attempt at providing rigorous mathematical analysis of preprocessing algorithms. While the initial interest in kernelization was driven mainly by practical applications, it turns out that kernelization provides a deep insight into the nature of fixed-parameter tractability. In the last few years, kernelization has transformed into a major research domain of Parameterized Complexity and many important advances in the area are on kernelization. These advances include general algorithmic findings on problems admitting kernels of

---

* This work was supported by the Netherlands Organization for Scientific Research (NWO), project "KERNELS: Combinatorial Analysis of Data Reduction".

polynomial size [1,4,15,21] and frameworks for ruling out polynomial kernels under certain complexity-theoretic assumptions [3,5,12,16].

A recent trend in the development of Parameterized Complexity, and more generally, Multivariate Analysis [23], is the study of the contribution of various secondary measurements (i.e., different than just the total input size or solution size) to problem complexity. Not surprisingly, the development of kernelization followed this trend resulting in various kernelization algorithms and complexity lower bounds for different kinds of parameterizations. In parameterized graph algorithms, one of the most important and relevant *complexity measures* of the graph is its treewidth. The algorithmic properties of problems parameterized by treewidth are, by now, well-understood. However, from the perspective of kernelization, this complexity measure is too general to obtain positive results: it is known for a multitude of graph problems such as VERTEX COVER, DOMINATING SET, and 3-COLORING, that there are no polynomial kernels parameterized by the treewidth of the input graphs unless NP $\subseteq$ coNP/poly [3]. This is why parameterization by more restrictive complexity measures, like the minimum size of a feedback vertex set or a vertex cover, is much more fruitful for kernelization.

In particular, kernelization of graph problems parameterized by the *vertex cover number*, which is the size of the smallest vertex set meeting all edges, was studied intensively [5,7,10,14,19]. For example, it has been shown that several graph problems such as VERTEX COVER, TREEWIDTH, and 3-COLORING, admit polynomial kernels parameterized by the size of a given vertex cover. On the other hand, under certain complexity-theoretic assumptions it is possible to show that a number of problems including DOMINATING SET [14], CLIQUE [5], and CHROMATIC NUMBER [5], do not admit polynomial kernels for this parameter. While different kernelization algorithms for various problems parameterized by vertex cover are known, we lack general a characterization of such problems. The main motivation of our work on this paper is the quest for meta-theorems on kernelization algorithms for problems parameterized by vertex cover.

According to Grohe [18], meta-theorems expose the deep relations between logic and combinatorial structures, which is a fundamental issue of computational complexity. Such theorems also yield a better understanding of the scope of general algorithmic techniques and the limits of tractability. The canonical example here is Courcelle's Theorem [9] which states that all problems expressible in Monadic Second-Order Logic are linear-time solvable on graphs of bounded treewidth. For more restricted parameters such as the vertex cover number, meta-theorems are available with a better dependency on the parameter [22,17]. In kernelization there are meta-theorems showing polynomial kernels for restricted graph families [4,15]. A systematic way to understand the kernelization complexity of parameterizations by vertex cover would therefore be to obtain a meta-theorem capturing a large class of problems admitting polynomial kernels. But is there a logic capturing the known positive results we are interested in? If such a logic exist, it would have to be able to express VERTEX COVER, which admits polynomial kernel, but not CLIQUE, which does not [20]; it should capture ODD CYCLE TRANSVERSAL and LONG CYCLE [20] but not DOMINAT-

ING SET [14]; and TREEWIDTH [7] but not CUTWIDTH [10]. As a consequence, if a logic capturing the phenomenon of polynomial kernelizability for problems parameterized by vertex cover exists, it should be a very strange logic and we therefore take a different approach.

In this paper, we provide two theorems with general conditions capturing a wide variety of known kernelization results about vertex cover parameterization. It has been observed before that reduction rules which identify irrelevant vertices by marking a polynomial number of vertices for each constant-sized subset of the vertex cover, lead to a polynomial kernel for several problems [6,19]. Our first contribution here is to uncover a characteristic of graph problems which explains their amenability to such reduction strategies, and to provide theorems using this characteristic. Roughly speaking, the problem of finding a minimum-size set of vertices which hits all induced subgraphs belonging to some family Π has a polynomial kernel parameterized by vertex cover, if membership in Π is invariant under changing the presence of all but a constant number of (non)edges incident to each vertex (and some technical conditions are met). The problem of finding the largest induced subgraph belonging to Π has a polynomial kernel parameterized by vertex cover under similar conditions. Our general theorems not only capture a wide variety of known results, they also imply results which were not known before. For example, as a corollary to our theorems we establish that the $\mathcal{F}$-MINOR-FREE DELETION deletion problem (i.e., for a fixed, finite list $\mathcal{F}$ of nonempty graphs, can we delete k vertices from G to ensure that the remaining graph does not contain a graph from $\mathcal{F}$ as a minor?) has a polynomial kernel for every fixed $\mathcal{F}$, when parameterized by the size of a vertex cover.

After studying the kernelization complexity of vertex-deletion and largest induced subgraph problems, we turn to two basic cases of property Π: containing some graph as an induced subgraph or minor. It is known that testing for a clique as an induced subgraph (when the desired size of the subgraph is part of the input) does not admit a polynomial kernel parameterized by vertex cover unless NP $\subseteq$ coNP/poly [5]. This is why we find the following result surprising: testing for a clique as a minor admits a polynomial kernel under the chosen parameterization. Driven by our desire to obtain a better understanding of the kernelization complexity of graph problems parameterized by vertex cover, we investigate induced subgraph testing and minor testing for other classes of graphs such as cycles, paths, matchings and stars. It turns out that the kernelization complexity of induced subgraph testing and minor testing is exactly opposite for all these classes. For example, testing for a star minor does not have a polynomial kernel due to its equivalence to CONNECTED DOMINATING SET [14], but we provide a polynomial kernel for testing the existence of an induced star subgraph by using a guessing step to reduce it to cases which are covered by our general theorems.

The paper is organized as follows. In Section 3 we describe a general reduction scheme, study its properties and use it to derive sufficient conditions for vertex-deletion and largest induced subgraph problems to admit polynomial kernels, parameterized by vertex cover. In Section 4 we investigate the kernelization complexity of induced subgraph versus minor testing for various graph families. A succinct overview of our results is given in Tables 1, 2, and 3.

2 Preliminaries

Parameterized complexity and kernels. A parameterized problem Q is a subset of $\Sigma^* \times \mathbb{N}$, the second component being the *parameter* which expresses some structural measure of the input. A parameterized problem is (strongly uniformly) *fixed-parameter tractable* if there exists an algorithm to decide whether $(x, k) \in Q$ in time $f(k)|x|^{\mathcal{O}(1)}$ where f is a computable function. We refer to the textbooks by Downey and Fellows, Flum and Grohe, and Niedermeier, for more background on parameterized complexity.

A *kernelization algorithm* (or *kernel*) for a parameterized problem Q is a polynomial-time algorithm which transforms an instance (x, k) into an equivalent instance (x', k') such that $|x'|, k' \leq f(k)$ for some computable function f, which is the *size* of the kernel. If $f \in k^{\mathcal{O}(1)}$ then this is a *polynomial kernel* (cf. [2]).

Graphs. All graphs we consider are finite, simple, and undirected. An undirected graph G consists of a vertex set $V(G)$ and a set of edges $E(G)$. A graph property Π is a (possibly infinite) set of graphs. The maximum degree of a vertex in G is denoted by $\Delta(G)$. A graph G is *empty* if $E(G) = \emptyset$. A vertex v is *simplicial* in graph G if $N_G(v)$ is a clique. A *minor model* of a graph H in a graph G is a mapping ϕ from $V(H)$ to subsets of $V(G)$ (called *branch sets*) which satisfies the following conditions: (a) $\phi(u) \cap \phi(v) = \emptyset$ for distinct $u, v \in V(H)$, (b) $G[\phi(v)]$ is connected for $v \in V(H)$, and (c) there is an edge between a vertex in $\phi(u)$ and a vertex in $\phi(v)$ for all $uv \in E(H)$. An *H-packing* in G is a set of vertex-disjoint subgraphs of G, each of which is isomorphic to H. An H-packing is *perfect* if the subgraphs cover the entire vertex set. The minimum size of a vertex cover in a graph G is denoted by $\textsc{vc}(G)$. For a set of vertices X in a graph G we use $G - X$ to denote the graph which results after deleting all vertices of X and their incident edges. We use the terms K_t and P_t to denote a clique or path on t vertices, respectively, whereas $K_{s,t}$ is a *biclique* (complete bipartite graph) whose partite sets have sizes s and t. The disjoint union of t copies of a graph G is represented by $t \cdot G$. Graph-theoretic concepts not defined here can be found in Diestel [13]. The set $\{1, 2, \ldots, n\}$ is abbreviated as $[n]$. If X is a finite set then $\binom{X}{n}$ denotes the collection of all subsets of X which have size exactly n. Similarly we use $\binom{X}{\leq n}$ for the subsets of size at most n (including $\emptyset$). The following proposition will be useful in showing that $\mathcal{F}$-Minor-Free Deletion can be captured by our general theorems.

Proposition 1. *If G contains H as a minor, then there is a subgraph $G^* \subseteq G$ containing an H-minor such that $\Delta(G^*) \leq \Delta(H)$ and $|V(G^*)| \leq |V(H)| + \textsc{vc}(G^*) \cdot (\Delta(H) + 1)$.*

Algorithm 1 Reduce(Graph G, Vertex cover $X \subseteq V(G), \ell \in \mathbb{N}, c_\Pi \in \mathbb{N}$)

```
for each $Y \in \binom{X}{\leq c_\Pi}$ and partition of $Y$ into $Y^+ \dot\cup Y^-$ do
    let $Z$ be the vertices in $V(G) \setminus X$ adjacent to all of $Y^+$ and to none of $Y^-$
    mark $\ell$ arbitrary vertices from $Z$ (if $|Z| < \ell$ then mark all of them)
delete from $G$ all unmarked vertices which are not contained in $X$
```

3 General Kernelization Theorems

3.1 Characterization by Few Adjacencies

In this section we introduce a general reduction rule for problems parameterized by vertex cover, and show that the rule preserves the existence of certain kinds of induced subgraphs. The central concept is the following.

Definition 1. *A graph property Π is* characterized by $c_\Pi \in \mathbb{N}$ adjacencies *if for all graphs $G \in \Pi$, for every vertex $v \in V(G)$, there is a set $D \subseteq V(G) \setminus \{v\}$ of size at most c_Π such that all graphs G' which are obtained from G by adding or removing edges between v and vertices in $V(G) \setminus D$, are also contained in Π.*

As an example of a property characterized by few adjacencies, consider the Hamiltonian graphs, i.e., the graphs which have a Hamiltonian cycle. This property is characterized by two adjacencies: given a graph G with a Hamiltonian cycle C and a vertex v, it is easy to see that as long as we preserve the edges between v and its predecessor and successor on C, changing the adjacency between v and other vertices preserves the Hamiltonicity of G. There are numerous other graph properties which are characterized by few adjacencies.

Proposition 2. *The following properties are characterized by constantly many adjacencies: (for any fixed finite set $\mathcal{F}$, graph H, or $\ell \geq 4$, respectively)*

1. *Containing $H \in \mathcal{F}$ as a minor ($c_\Pi = \max_{H \in \mathcal{F}} \Delta(H)$).*
2. *Having a perfect H-packing ($c_\Pi = \Delta(H)$).*
3. *Having a chordless cycle of length at least ℓ ($c_\Pi = \ell - 1$).*
4. *Having a Hamiltonian path (resp. cycle), or having an odd cycle ($c_\Pi = 2$).*

As an illustrative non-example, note that the properties of being a cycle, of having chromatic number at least four, or of *not* being a perfect graph, cannot be characterized by a constant number of adjacencies.

The single reduction rule that we will use to derive our general kernelization theorems, is the Reduce procedure presented as Algorithm 1. Its utility for kernelization stems from the fact that it efficiently shrinks a graph to a size bounded polynomially in the cardinality of the given vertex cover X.

Observation 1 *For every fixed constant c_Π,* Reduce(G, X, ℓ, c_Π) *runs in polynomial time and results in a graph on $\mathcal{O}(|X| + \ell \cdot 2^{c_\Pi} \cdot |\binom{X}{\leq c_\Pi}|) = \mathcal{O}(|X| + \ell \cdot |X|^{c_\Pi})$ vertices.*

The soundness of the Reduce procedure for many types of kernelization comes from the following lemma. It shows that for graph properties Π which are characterized by few adjacencies, an application of Reduce with parameter $\ell = s + p$ preserves the existence of induced Π subgraphs of size up to p that avoid any set of size at most s.

Lemma 1. *Let Π be characterized by c_Π adjacencies, and let G be a graph with vertex cover X. If $G[P] \in \Pi$ for some $P \subseteq V(G) \setminus S$ and $S \subseteq V(G)$, then for any $\ell \geq |S| + |P|$ the graph G' resulting from* Reduce(G, X, ℓ, c_Π) *contains $P' \subseteq V(G') \setminus S$ such that $G'[P'] \in \Pi$ and $|P'| = |P|$.*

3.2 Kernelization for Vertex-Deletion Problems

We will present a general theorem which gives polynomial kernels for vertex-deletion problems of the following form.

Deletion Distance To Π-free (vc)
Input: A graph G with vertex cover $X \subseteq V(G)$, and an integer $k \geq 1$.
Parameter: The size $|X|$ of the vertex cover.
Question: Is there a set $S \subseteq V(G)$ of size at most k such that $G - S$ does not contain a graph in Π as an induced subgraph?

Observe that Π need not be finite or decidable. The condition that a vertex cover is given along with the input is present for technical reasons; to apply the data reduction schemes presented in this paper, one may simply compute a 2-approximate vertex cover and use that as X.

Theorem 1. *If Π is a graph property such that:*

(i) Π is characterized by c_Π adjacencies,
(ii) every graph in Π contains at least one edge, and
(iii) there is a non-negative polynomial $p\colon \mathbb{N} \to \mathbb{N}$ such that all graphs $G \in \Pi$ contain an induced subgraph $G' \in \Pi$ such that $V(G') \leq p(\mathrm{vc}(G'))$,

then Deletion Distance To Π-free (vc) *has a kernel with $\mathcal{O}((x+p(x))x^{c_\Pi})$ vertices, where $x := |X|$.*

Before proving the theorem, let us briefly discuss its preconditions. We cannot drop Property (ii), as otherwise the theorem would capture the Clique problem (taking $\Pi := \{2 \cdot K_1\}$), for which a lower bound exists [5]. If Property (i) is dropped, then the theorem would capture problems such as Perfect Deletion for which the kernelization complexity is still open. We require the third condition to make the proof go through.

Proof (of Theorem 1). Consider some input instance (G, X, k). Firstly, observe that if $k \geq |X|$, then we clearly have a yes-instance: removal of X results in an edgeless graph, which is guaranteed not to contain induced subgraphs from Π

Table 1. Problems which admit polynomial kernels when parameterized by the size of a given vertex cover, by applying Theorem 1. For $\mathcal{F}$-MINOR-FREE DELETION we require that no graph in $\mathcal{F}$ is empty (cf. Section E).

Problem	Π	c_{Π}
VERTEX COVER	$\{K_2\}$	1
ODD CYCLE TRANSVERSAL	Graphs containing an odd cycle	2
CHORDAL DELETION	Graphs with a chordless cycle	3
PLANARIZATION	Graphs with a K_5 or $K_{3,3}$ minor	4
η-TRANSVERSAL (cf. [11])	Graphs of treewidth $> \eta$	$f(\eta)$
$\mathcal{F}$-MINOR-FREE DELETION	Graphs with an $H \in \mathcal{F}$-minor	$\max_{H \in \mathcal{F}} \Delta(H)$

due to Property (ii). Therefore, we may assume that $k < |X|$ as otherwise we output a trivial YES-instance.

We let G' be the result of REDUCE$(G, X, k + p(|X|), c_{\Pi})$ and return the instance (G', X, k), which gives the right running time and size bound by Observation 1. We need to prove that the output instance (G', X, k) is equivalent to the input instance (G, X, k). As G' is an induced subgraph of G, it follows that if $G - S$ does not contain any graph in Π, then neither does $G' - (S \cap V(G'))$. Therefore, if (G, X, k) is a YES-instance, then so is (G', X, k). Assume then, that (G', X, k) is a YES-instance and let S be a subset of vertices with $|S| \leq k$ such that $G' - S$ does not contain any induced subgraph from Π. We claim that $G - S$ does not contain such induced subgraphs either, i.e., that S is also a feasible solution for the instance (G, X, k).

Assume for the sake of contradiction that there is a set $P \subseteq V(G) \setminus S$ such that $G[P] \in \Pi$. Consider a *minimal* such set P, which ensures by Property (iii) that $|P| \leq p(\text{VC}(G[P]))$. As $P \cap X$ is a vertex cover of $G[P]$, it follows that $|P| \leq p(|P \cap X|) \leq p(|X|)$. As we executed the reduction with parameter $\ell = k + p(|X|)$, Lemma 1 guarantees the existence of a set $P' \subseteq V(G') \setminus S$ such that $G'[P'] \in \Pi$. But this shows that the graph $G' - S$ contains an induced Π subgraph, contradicting the assumption that S is a solution for G' and thereby concluding the proof. □

Corollary 1. *All problems in Table 1 fit into the framework of Theorem 1 and admit polynomial kernels parameterized by the size of a given vertex cover.*

Proof. We give the proof for $\mathcal{F}$-MINOR-FREE DELETION; the proofs for the other items can be found in Section C.3. If we let Π contain all graphs that contain a member of $\mathcal{F}$ as a minor, then a graph is Π-induced-subgraph-free if and only if it is $\mathcal{F}$-minor-free. By Proposition 2 this class Π is characterized by $c_{\Pi} := \max_{H \in \mathcal{F}} \Delta(H)$ adjacencies, so we satisfy Property (i). Since we demand all members of $\mathcal{F}$ to be nonempty, the graphs containing a minor from $\mathcal{F}$ have at least one edge (Property (ii)). Finally, consider a vertex-minimal graph G^* which contains a graph $H \in \mathcal{F}$ as a minor. By Proposition 1 we have $|V(G^*)| \leq |V(H)| + \text{VC}(G^*) \cdot (\Delta(H) + 1)$. As $\mathcal{F}$ is fixed, the maximum degree and size of graphs in $\mathcal{F}$ are constants which shows that Property (iii) is satisfied, resulting in a kernel with $\mathcal{O}(|X|^{\Delta+1})$ vertices for $\Delta := \max_{H \in \mathcal{F}} \Delta(H)$. □

Table 2. Problems which admit polynomial kernels when parameterized by the size of a given vertex cover, by applying Theorem 2 (as proven in Corollary 2).

Problem	Π	c_Π
LONG CYCLE	Graphs with a Hamiltonian cycle	2
LONG PATH	Graphs with a Hamiltonian path	2
H-PACKING for nonempty H	Graphs with a perfect H-packing	$\Delta(H)$

3.3 Kernelization for Largest Induced Subgraph Problems

In this section we study the following class of problems, which is in some sense dual to the class considered previously.

> LARGEST INDUCED Π-SUBGRAPH (VC)
> **Input:** A graph G with vertex cover $X \subseteq V(G)$, and an integer $k \geq 1$.
> **Parameter:** The size $|X|$ of the vertex cover.
> **Question:** Is there a set $P \subseteq V(G)$ of size at least k such that $G[P] \in \Pi$?

The following theorem gives sufficient conditions for the existence of polynomial kernels for such problems.

Theorem 2. *If Π is a graph property such that:*

- *(i) Π is characterized by c_Π adjacencies, and*
- *(ii) there is a non-negative polynomial $p\colon \mathbb{N} \to \mathbb{N}$ such that for all graphs $G \in \Pi$, $|V(G)| \leq p(\textsc{vc}(G))$,*

then LARGEST INDUCED Π-SUBGRAPH (VC) *has a kernel with $\mathcal{O}(p(|X|)|X|^{c_\Pi})$ vertices.*

The proof is similar to that of Theorem 1, and can be found in the appendix.

4 Subgraph Testing versus Minor Testing

Several important graph problems such as CLIQUE, LONG PATH, and LONG INDUCED PATH, can be stated in terms of testing for the existence of a certain graph H as an induced subgraph, or as a minor. Note that for these problems, the size of the graph whose containment in G is tested is part of the input as the problem is polynomial-time solvable for each constant size. We compared the kernelization complexity of induced subgraph- versus minor testing for various types of graphs, parameterized by vertex cover, and found the surprising outcome that the kernelization complexity is often opposite: one variant admits a polynomial kernel while the other does not, assuming NP $\not\subseteq$ coNP/poly. We discuss our findings separately for each type of graph whose containment is tested.

Table 3. Kernelization complexity of testing for induced H subgraphs versus testing for H as a minor, when the graph H is given as part of the input by specifying t. The problems are parameterized by the size of a given vertex cover. Kernel lower bounds are under the assumption that NP $\not\subseteq$ coNP/poly.

Graph H	Testing for induced H	Testing for H-minor
K_t	$\neg\, \exists \lvert X\rvert^{\mathcal{O}(1)}$ kernel [5]	$\exists \lvert X\rvert^{\mathcal{O}(1)}$ kernel (Thm. 3)
$K_{1,t}$	$\exists \lvert X\rvert^{\mathcal{O}(1)}$ kernel (Thm. 5)	$\neg\, \exists \lvert X\rvert^{\mathcal{O}(1)}$ kernel [14]
$K_{s,t}$	$\neg\, \exists \lvert X\rvert^{\mathcal{O}(1)}$ kernel (Thm. 6)	$\neg\, \exists \lvert X\rvert^{\mathcal{O}(1)}$ kernel [14]
P_t	$\neg\, \exists \lvert X\rvert^{\mathcal{O}(1)}$ kernel (Thm. 7)	$\exists \lvert X\rvert^{\mathcal{O}(1)}$ kernel ([6] or Thm. 2)
$t \cdot K_2$	$\neg\, \exists \lvert X\rvert^{\mathcal{O}(1)}$ kernel (Thm. 8)	P-time solvable

4.1 Testing for Cliques

The Clique problem (i.e., testing for K_t as an induced subgraph) was one of the first problems known not to admit a polynomial kernel parameterized by the size of a given vertex cover [5, Theorem 11]. Our main result of this section is a polynomial kernel for the related minor testing problem.

> Clique Minor Test (vc)
> **Input:** A graph G with vertex cover $X \subseteq V(G)$, and an integer $t \geq 1$.
> **Parameter:** The size $|X|$ of the vertex cover.
> **Question:** Does G contain K_t as a minor?

All problems we study in Section 4 are defined similarly in the obvious way, and we will not define them explicitly. Our polynomial kernel uses reduction rules based on simplicial vertices, inspired by the recent work on kernels for Treewidth [7].

Theorem 3. Clique Minor Test (vc) *admits a kernel with* $\mathcal{O}(|X|^4)$ *vertices.*

Firstly, observe that if a graph has a clique K_t as a minor, then its vertex cover number is at least $t-1$: taking a minor does not increase the vertex cover number, and $\textsc{vc}(K_t) = t-1$. Therefore, we assume that $t \leq |X|+1$, as otherwise we may output a trivial no-instance. Our algorithm is based on three reduction rules. In the following, we assume the reduction rules are exhaustively applied in their given order.

Reduction Rule 1 *For every distinct pair* $v, w \in X$ *such that* $vw \notin E(G)$*, if there are more than* $(|X|+1)^2$ *vertices in* $V(G) \setminus X$ *adjacent both to* v *and* w*, then add the edge* vw*. Output the resulting instance* (G', X, t)*.*

Lemma 2. *Rule 1 is safe.*

Proof. As G is a subgraph of G', any clique minor in G is also contained in G'. Therefore we need to argue that if G' admits a K_t minor, then so does G.

Assume that G' has a K_t minor, and let G^* be a subgraph of G' containing a K_t minor such that $|V(G^*)| \leq |V(K_t)| + \textsc{vc}(G') \cdot (\Delta(K_t)+1) = t + \textsc{vc}(G') \cdot t$, whose existence is guaranteed by Proposition 1. As $\text{VC}(G') \leq |X|$ it follows that $|\bigcup_{v \in K_t} \phi(v)| \leq t + |X| \cdot t$, and since $t \leq |X| + 1$ the number of vertices involved in the minor model is at most $(|X| + 1)^2$. Hence by the precondition to the reduction rule, there is a vertex y adjacent to both v and w which is not used in the minor model.

Observe that if ϕ avoids one of v and w, it is also a clique model in G. Assume then that $v \in \phi(u_1)$ and $w \in \phi(u_2)$; it may happen that $u_1 = u_2$. Now we can transform ϕ into a clique minor model ϕ' in G, by adding y to $\phi(u_1)$: contraction of the edge vy in this branch set creates the edge vw that was missing in G. $\square$

Reduction Rule 2 *If there exists a simplicial vertex $s \in V(G) \setminus X$ such that $\deg(s) \geq t - 1$, output a trivial* YES*-instance.*

Reduction Rule 3 *If there exists a simplicial vertex $s \in V(G) \setminus X$ such that $\deg(s) < t - 1$, delete it. Output the resulting instance (G', X, t).*

Correctness of Rule 2 is obvious, as s together with its neighborhood already forms a K_t. The correctness proof for Rule 3 can be found in the appendix (Lemma 3). The running time of the kernelization algorithm is polynomial, as the presented reduction rules can only add edges inside X and remove vertices from $V(G) \setminus X$. Exhaustive application of the reduction rules results in an instance with at most $(|X| + 1)^4$ vertices (Lemma 4), which proves Theorem 3.

4.2 Testing for Bicliques

We now consider testing for a biclique as an induced subgraph or minor. Observe first that if G is a connected graph on at least three vertices, then the following conditions are equivalent: graph G has a (a) spanning tree with t or more leaves, (b) $K_{1,t}$ minor, (c) connected dominating set of size at most $|V(G)| - t$. Hence there is a trivial polynomial-parameter transformation [2] from Connected Dominating Set (vc) to $K_{1,t}$ Minor Test (vc). Dom et al. [14, Theorem 5] showed[3] that the former problem does not admit polynomial kernels unless NP $\subseteq$ coNP/poly, and hence the same lower bound holds for the latter.

The situation is more diverse when testing for a biclique as an induced subgraph. If we fix a constant c and wish to test for a biclique $K_{c,t}$ as induced subgraph, where t is part of the input, then this problem admits a polynomial kernel parameterized by vertex cover. The kernel is given as Theorem 5 in the appendix. Our main insight is a polynomial-size compression which is obtained by guessing the model of the constant-size partite set within the vertex cover, reducing the problem to the OR of $\binom{|X|}{c}$ instances of Independent Set parameterized by vertex cover. As Independent Set parameterized by vertex cover is equivalent to Vertex Cover parameterized by the size of a given (suboptimal) vertex cover, each of these can be compressed to a size polynomial in $|X|$ using

---

[3] The lower bound they give is for Dominating Set, but a trivial transformation extends it to Connected Dominating Set.

Theorem 1, and the NP-completeness transformation then results in an instance of the original problem of size $\mathcal{O}(|X|)^{\mathcal{O}(1)}$ which forms the kernel.

If the sizes of both partite sets are part of the input, then we can no longer obtain a polynomial kernel. Theorem 6 in the appendix gives a cross-composition [5] from Balanced Biclique in Bipartite Graphs to show that testing for an induced $K_{s,t}$ subgraph, parameterized by vertex cover, does not admit a polynomial kernel unless NP $\subseteq$ coNP/poly.

4.3 Testing for Paths

Since a graph contains P_t as a minor if and only if it contains P_t as a subgraph, testing for a P_t minor is equivalent to the Long Path problem and hence has a polynomial kernel parameterized by vertex cover, through Theorem 2. The related induced subgraph testing problem, however, is unlikely to admit a polynomial kernel. We cross-compose t instances of Hamiltonian $s-t$ Path into a single instance of Long Induced Path (vc). The main idea behind the construction is to create an instance containing three paths P_A, P_B, P_C of consecutive degree-two vertices, such that any sufficiently long induced path traverses all these paths. The only connections between P_A and P_B can be made by visiting a vertex v_i outside the vertex cover; there is one such vertex v_i for each input instance. To make a suitably long path, a solution must traverse P_A, then visit some v_i, and then continue traversing P_B. The inducedness requirement ensures the path cannot visit neighbors of v_i except for its predecessor on P_A and successor on P_B. This allows us to encode the adjacency matrix of the input graph corresponding to v_i into the set of edges incident on v_i. The full proof is given as Theorem 7, and shows that Long Induced Path (vc) does not admit a polynomial kernel unless NP $\subseteq$ coNP/poly.

4.4 Testing for Matchings

Matchings (i.e., disjoint unions of K_2's) are the last type of graphs whose containment testing we consider. It is not difficult to see that G has a $t \cdot K_2$ minor if and only if G has a matching of size t, and hence we can solve the minor-testing variant of this containment problem in polynomial time by simply computing a maximum matching. On the other hand, finding an induced matching is a classic NP-complete problem and we give evidence that it does not admit a polynomial kernel parameterized by vertex cover. We use a bit-selector strategy to cross-compose Maximum Induced Matching in Bipartite Graphs into our target problem, exploiting the inducedness requirement to allow the bitselector to isolate a solution corresponding to a single input instance. Hence we prove that Maximum Induced Matching (vc) does not admit a polynomial kernel unless NP $\subseteq$ coNP/poly (Theorem 8 in the appendix).

5 Conclusion

We have studied the existence of polynomial kernels for graph problems parameterized by vertex cover. The general theorems we presented unify known positive results for many problems, and the characterization in terms of forbidden or desired induced subgraphs from a class characterized by few adjacencies gives a common explanation for the results obtained earlier. Our comparison of induced subgraph and minor testing problems shows that the kernelization complexity landscape of problems parameterized by vertex cover is rich and difficult to capture with a single meta-theorem. The kernel lower bounds for INDUCED $K_{s,t}$ SUBGRAPH TEST (VC), LONG INDUCED PATH (VC), and MAXIMUM INDUCED MATCHING (VC), show that besides connectivity and domination requirements, an inducedness requirement can form an obstacle to kernelizability for parameterizations by vertex cover.

An obvious direction for further work is to find even more general kernelization theorems which can also encompass the known positive results for problems like TREEWIDTH (VC) [7] and CLIQUE MINOR TEST (VC). There are also various problems for which the kernelization complexity parameterized by vertex cover is still open; among these are PERFECT DELETION, INTERVAL DELETION, BANDWIDTH and GENUS. Finally, one may investigate whether Theorem 1 has an analogue for edge-deletion problems.

References

1. N. Alon, G. Gutin, E. J. Kim, S. Szeider, and A. Yeo. Solving MAX-r-SAT above a tight lower bound. *Algorithmica*, 61(3):638–655, 2011.
2. H. L. Bodlaender. Kernelization: New upper and lower bound techniques. In *Proc. 4th IWPEC*, pages 17–37, 2009. `doi:10.1007/978-3-642-11269-0_2`.
3. H. L. Bodlaender, R. G. Downey, M. R. Fellows, and D. Hermelin. On problems without polynomial kernels. *J. Comput. Syst. Sci.*, 75(8):423–434, 2009. `doi:10.1016/j.jcss.2009.04.001`.
4. H. L. Bodlaender, F. V. Fomin, D. Lokshtanov, E. Penninkx, S. Saurabh, and D. M. Thilikos. (Meta) Kernelization. In *Proc. 50th FOCS*, pages 629–638, 2009. `doi:10.1109/FOCS.2009.46`.
5. H. L. Bodlaender, B. M. P. Jansen, and S. Kratsch. Cross-composition: A new technique for kernelization lower bounds. In *Proc. 28th STACS*, pages 165–176, 2011. `doi:10.4230/LIPIcs.STACS.2011.165`.
6. H. L. Bodlaender, B. M. P. Jansen, and S. Kratsch. Kernel bounds for path and cycle problems. In *Proc. 6th IPEC*, pages 145–158, 2011. `doi:10.1007/978-3-642-28050-4_12`.
7. H. L. Bodlaender, B. M. P. Jansen, and S. Kratsch. Preprocessing for treewidth: A combinatorial analysis through kernelization. In *Proc. 38th ICALP*, pages 437–448, 2011. `doi:10.1007/978-3-642-22006-7_37`.
8. K. Cameron. Induced matchings. *Discrete Applied Mathematics*, 24(1-3):97–102, 1989. `doi:10.1016/0166-218X(92)90275-F`.
9. B. Courcelle. The monadic second-order logic of graphs I: Recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990. `doi:10.1016/0890-5401(90)90043-H`.

10. M. Cygan, D. Lokshtanov, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. On cutwidth parameterized by vertex cover. In *Proc. 6th IPEC*, pages 246–258, 2011. `doi:10.1007/978-3-642-28050-4_20`.
11. M. Cygan, D. Lokshtanov, M. Pilipczuk, M. Pilipczuk, and S. Saurabh. On the hardness of losing width. In *Proc. 6th IPEC*, pages 159–168, 2011. `doi:10.1007/978-3-642-28050-4_13`.
12. H. Dell and D. van Melkebeek. Satisfiability allows no nontrivial sparsification unless the polynomial-time hierarchy collapses. In *Proc. 42nd STOC*, pages 251–260, 2010. `doi:10.1145/1806689.1806725`.
13. R. Diestel. *Graph Theory*. Springer-Verlag, Heidelberg, 4th edition, 2010.
14. M. Dom, D. Lokshtanov, and S. Saurabh. Incompressibility through colors and IDs. In *Proc. 36th ICALP*, pages 378–389, 2009. `doi:10.1007/978-3-642-02927-1_32`.
15. F. Fomin, D. Lokshtanov, S. Saurabh, and D. M. Thilikos. Bidimensionality and kernels. In *Proc. 21st SODA*, pages 503–510, 2010.
16. L. Fortnow and R. Santhanam. Infeasibility of instance compression and succinct PCPs for NP. *J. Comput. Syst. Sci.*, 77(1):91–106, 2011. `doi:10.1016/j.jcss.2010.06.007`.
17. R. Ganian. Twin-cover: Beyond vertex cover in parameterized algorithmics. In *Proc. 6th IPEC*, pages 259–271, 2011. `doi:10.1007/978-3-642-28050-4_21`.
18. M. Grohe. *Logic, Graphs, and Algorithms*, chapter in Logic and Automata: History and Perspectives, pages 357 – 422. Amsterdam University Press, 2007.
19. B. M. P. Jansen and S. Kratsch. Data reduction for graph coloring problems. In *Proc. 18th FCT*, pages 90–101, 2011. `doi:10.1007/978-3-642-22953-4_8`.
20. B. M. P. Jansen and S. Kratsch. On polynomial kernels for structural parameterizations of odd cycle transversal. In *Proc. 6th IPEC*, pages 132–144, 2011. `doi:10.1007/978-3-642-28050-4_11`.
21. S. Kratsch and M. Wahlström. Compression via matroids: a randomized polynomial kernel for odd cycle transversal. In *Proc. 23rd SODA*, pages 94–103, 2012.
22. M. Lampis. Algorithmic meta-theorems for restrictions of treewidth. In *Proc. 18th ESA*, pages 549–560, 2010. `doi:10.1007/978-3-642-15775-2_47`.
23. R. Niedermeier. Reflections on multivariate algorithmics and problem parameterization. In *Proc. 27th STACS*, pages 17–32, 2010. `doi:10.4230/LIPIcs.STACS.2010.2495`.

A Preliminaries for Kernelization Lower Bounds

To prove lower bounds we frequently use the framework of cross-composition [5], which builds on earlier work by Bodlaender et al. [3], and Fortnow and Santhanam [16].

Definition 2 (Polynomial equivalence relation [5]). *An equivalence relation $\mathcal{R}$ on Σ^* is called a* polynomial equivalence relation *if the following two conditions hold:*

1. *There is an algorithm that given two strings $x, y \in \Sigma^*$ decides whether x and y belong to the same equivalence class in $(|x| + |y|)^{\mathcal{O}(1)}$ time.*
2. *For any finite set $S \subseteq \Sigma^*$ the equivalence relation $\mathcal{R}$ partitions the elements of S into at most $(\max_{x \in S} |x|)^{\mathcal{O}(1)}$ classes.*

Definition 3 (Cross-composition [5]). *Let $L \subseteq \Sigma^*$ be a set and let $Q \subseteq \Sigma^* \times \mathbb{N}$ be a parameterized problem. We say that L* cross-composes *into Q if there is a polynomial equivalence relation $\mathcal{R}$ and an algorithm which, given r strings $x_1, x_2, \ldots, x_r$ belonging to the same equivalence class of $\mathcal{R}$, computes an instance $(x^*, k^*) \in \Sigma^* \times \mathbb{N}$ in time polynomial in $\sum_{i=1}^{r} |x_i|$ such that:*

1. *$(x^*, k^*) \in Q \Leftrightarrow x_i \in L$ for some $1 \leq i \leq r$,*
2. *k^* is bounded by a polynomial in $\max_{i=1}^{r} |x_i| + \log r$.*

Theorem 4 ([5]). *If some set $L \subseteq \Sigma^*$ is NP-hard under Karp reductions and L cross-composes into the parameterized problem Q then there is no polynomial kernel for Q unless NP $\subseteq$ coNP/poly.*

B Omitted proofs for Section 2

Proof (of Proposition 1). Let G be a graph containing a model ϕ of a graph H. We show how to find a subgraph G^* satisfying the claims.

For every pair of distinct vertices $u, v \in V(H)$, mark an arbitrary edge between $\phi(u)$ and $\phi(v)$. In each branch set $\phi(v)$ for $v \in V(H)$, mark the edges of any inclusion-minimal tree T_v that contains all the vertices incident to edges marked so far. For each isolated vertex $v \in V(H)$, mark an arbitrary vertex in $\phi(v)$. Now obtain G^* from G by deleting unmarked edges, and deleting unmarked vertices which are not incident on a marked edge. It is easy to verify that restricting ϕ to G^* gives an H-model in G^*. To see that $\Delta(G^*) \leq \Delta(H)$, consider a vertex $v \in V(G^*)$ and partition the edges incident on it into two types: those which were marked to build a tree T_u in a branch set $\phi(u)$ for some $u \in V(H)$, and those which connect two different branch sets. Suppose v is incident on ℓ edges of the tree T_u. Then T_u has at least ℓ leaves, and all these leaves connect $\phi(u)$ to different branch sets. Observe that each connection to a different branch set corresponds to a distinct neighbor of u in H. As u has at most $\Delta(H)$ neighbors in H, there are at most $\Delta(H)$ connections between the branch set $\phi(u)$ and other branch sets. Since at least ℓ connections are made by leaves of T_u

unequal to v, edges incident on v can make at most $\Delta(H) - \ell$ connections to other branch sets. As this accounts for all edges incident on v in G^* it follows that $\deg_{G^*}(v) \leq \ell + (\Delta(H) - \ell)$. As v was arbitrary this proves $\Delta(G^*) \leq \Delta(H)$. It remains to prove that $|V(G^*)|$ is suitably small.

Let $X \subseteq V(G^*)$ be a minimum vertex cover of G^*. All isolated vertices in G^* correspond to isolated vertices in H, so there are at most $|V(H)|$ of them. The remaining vertices of G^* which do not belong to X, have at least one neighbor in X (as X is a vertex cover and the vertices are not isolated). Since each vertex in X has degree at most $\Delta(H)$, the total number of vertices in G^* is at most $|V(H)| + |X| + \Delta(H) \cdot |X| \leq |V(H)| + |X|(\Delta(H) + 1)$ which proves the claim. □

C Omitted proofs for Section 3

C.1 Proof of Proposition 2

Proof (of Proposition 2). We prove the claims one by one.

(1) Let $\mathcal{F}$ be a finite set of graphs. The property of having an $H \in \mathcal{F}$-minor is characterized by $c_\Pi := \max_{H \in \mathcal{F}} \Delta(H)$ adjacencies. Let G contain $H' \in \mathcal{F}$ as a minor, and let $v \in V(G)$ be an arbitrary vertex. We give a set $D \subseteq V(G) \setminus \{v\}$ of size at most $\max_{H \in \mathcal{F}} \Delta(H)$ such that changing the adjacencies between v and $V(G) \setminus D$ preserves the fact that G has an H'-minor. By Proposition 1 a subgraph G^* of G with maximum degree at most $\Delta(H')$ exists, which has an H'-minor model ϕ. If v is not contained in graph G^*, then changing the presence of edges incident to v preserves the minor model ϕ in G. If v is contained in G^*, then pick $D := N_{G^*}(v)$ which has size at most $\Delta(H')$ by the degree bound of G^* guaranteed by the proposition. Changing adjacencies between v and $V(G) \setminus D$ preserves the fact that G^* is a subgraph of G, and therefore preserves the fact that G has H' as a minor; this implies membership in Π.

(2) We will prove the property is characterized by $c_\Pi := \Delta(H)$ adjacencies. Fix a graph H and let G be a graph with a perfect H-packing. For an arbitrary vertex $v \in V(G)$, consider a perfect H-packing in G and let G' be the subgraph in the packing which contains v. Picking $D := N_{G'}(v)$ it follows that $|D| \leq \Delta(H)$. Changing adjacencies between v and $V(G) \setminus D$ in G preserves the perfect H-packing we started from, as all edges incident on v needed to make the subgraph G' isomorphic to H are maintained. Hence the graph resulting from such modifications has a perfect H-packing and is contained in Π.

(3) We prove that the property of having a chordless cycle is characterized by $c_\Pi := \ell - 1$ adjacencies. Let G be a graph with a chordless cycle C of length at least ℓ, and let v be an arbitrary vertex. If v does not lie on C then changing the presence of edges incident on v preserves C and results in a graph with a chordless cycle of length at least ℓ. Suppose that v lies on C, and label the vertices on C as $(v, v_2, \ldots, v_k)$ for $k \geq \ell$. Define $D := \{v_1, v_2, \ldots, v_{\ell-1}, v_k\}$, i.e., D contains the predecessor of v and its $\ell - 2$ successors. Let G' be obtained from G by changing the adjacency between v and $V(G) \setminus D$. We prove that G' has a chordless cycle of

length at least ℓ. By definition of G' it follows that $G'[\{v, v_1, \ldots, v_{\ell-1}]$ is an induced path, as we did not change any adjacencies here. Since $(v_{\ell-1}, v_\ell, \ldots, v_k, v)$ was an induced path in G, and we explicitly preserve the edge between v_k and v, it follows that $(v_{\ell-1}, v_\ell, \ldots, v_k, v)$ is a path in G', but it might have chords in G. Take a shortest path from $v_{\ell-1}$ to v_k through the set $\{v_{\ell-1}, v_\ell, \ldots, v_k, v\}$; by the previous observations such a path exists, and since it is a shortest path, it is chordless. As the definition of G' ensures the adjacencies between v_i and v_j have not changed for any $i, j \in \{2, \ldots, k\}$ it follows that appending the chordless $v_{\ell-1}$ to v path onto the chordless path $(v, v_2, \ldots, v_{\ell-1}$ gives a chordless cycle in G' of length at least ℓ, completing the proof.

(4) First consider the property of being Hamiltonian. The property is characterized by $c_\Pi := 2$ adjacencies. Take a graph G with a Hamiltonian cycle. For any vertex v, fixing its adjacency to its predecessor and successor on the cycle is sufficient to preserve the existence of the cycle in any graph obtained by changing the presence of other edges incident on v. As the length of the cycle is not affected, the same proof goes for the property of having an odd cycle. The property of having a Hamiltonian path is similar; for the endpoints we only have to preserve a single adjacency. □

C.2 Proof of Lemma 1

Proof (of Lemma 1). Assume the conditions in the lemma statement hold, and let $R \subseteq V(G)$ be the vertices which are removed by the reduction procedure, i.e., $R := V(G) \setminus V(G')$. Let $p_1, p_2, \ldots, p_t$ be an arbitrary ordering of $P \cap R$. We inductively create a sequence of sets $P = P_0, P_1, \ldots, P_t$ such that (a) $G[P_i] \in \Pi$, (b) $|P_i| = |P|$, (c) $P_i \cap S = \emptyset$, and (d) $P_i \cap R = \{p_{i+1}, p_{i+2}, \ldots, p_t\}$ for every $i \in \{0, 1, \ldots, \ell\}$. Note that P satisfies the constraints imposed on P_0, while existence of P_ℓ proves the lemma. Hence, we only need to show how to construct P_i from P_{i-1} for $i \in \{1, 2, \ldots, t\}$.

Consider graph $G[P_{i-1}]$ and vertex $p_i \in P_{i-1}$. As $G[P_{i-1}] \in \Pi$, Definition 1 ensures that there exists a set D of at most c_Π vertices of P_{i-1} such that arbitrarily changing adjacencies between p_i and vertices of $P_{i-1} \setminus D$ in $G[P_{i-1}]$ preserves membership in Π. Let $D^+ := N_G(p_i) \cap D$ and $D^- := D \setminus D^+$. Since vertex p_i is contained in R and was removed by the reduction process, it follows from the deletion procedure that $p_i \notin X$ and therefore that $D^+ \subseteq N_G(p_i) \subseteq X$ since X is a vertex cover of G. Let $D^-_X := D^- \cap X$. Observe that p_i was a candidate for marking for the partition (D^+, D^-_X) of $D \cap X$, but as $p_i \in R$ it was not marked. Hence, there exist $\ell \geq |S| + |P|$ marked vertices in $V(G) \setminus X$ adjacent to all of D^+ and none of D^-_X. As $|P_{i-1}| = |P|$ and p_i is not marked, we can find a vertex $p'_i \in V(G) \setminus X$ that does not belong to P_{i-1} or S, is marked, and has the same neighborhood in $D \cap X$ as p_i. Since X is a vertex cover, both p_i and p'_i have all their neighbors in X. As p'_i is not adjacent to any member of D^-_X, it is not adjacent to D^-. Take $P_i := (P_{i-1} \cup \{p'_i\}) \setminus \{p_i\}$. Note that $|P_i| = |P_{i-1}| = |P|$ and $P_i \cap R = \{p_{i+1}, p_{i+2}, \ldots, p_t\}$. Moreover, a graph isomorphic to $G[P_i]$ can be obtained from $G[P_{i-1}]$ by changing adjacencies between p_i and vertices of $P_{i-1} \setminus D$. The only adjacencies that need to be changed are between p_i and $N_G(p_i) \triangle N_G(p'_i) \subseteq X$ ($\triangle$ denotes symmetric difference), but this set is disjoint with D and hence the changes preserve membership in Π. As P_i satisfies all induction claims, this completes the proof. $\square$

C.3 Proof of Corollary 1

We present the remainder of the proof.

Proof (of the remainder of Corollary 2). We consider the problems in the order of Table 1 and show how they fit into the framework.

VERTEX COVER (VC). Observe that a graph G has a vertex cover of size k if and only there is a set $S \subseteq V(G)$ such that $G - S$ is an independent set, or equivalently, $G - S$ does not have K_2 as an induced subgraph. So VERTEX COVER (VC) is equivalent to DELETION DISTANCE TO Π-FREE (VC) for $\Pi = \{K_2\}$. Since this Π contains only a single graph of degree one, it is easily seen to be characterized by the single adjacency of one vertex in K_2 to its neighbor (Property (i)). Obviously all graphs in Π contain at least one edge (Property (ii)), and since Π contains a single graph on two vertices, having a vertex cover of size one, the constant function $p(n) := 2$ suffices for Property (iii). Hence all preconditions for Theorem 1 are satisfied and the problem has a kernel with $\mathcal{O}(|X|^2)$ vertices.

ODD CYCLE TRANSVERSAL (VC). A graph G is bipartite if and only if it does not contain a graph with an odd cycle as an induced subgraph. Hence by letting Π contain all graphs which contain an odd cycle (which is not the same as letting Π be the class of all odd cycles), ODD CYCLE TRANSVERSAL (VC) is equivalent to DELETION DISTANCE TO Π-FREE (VC). By Proposition 2 this property Π is characterized by a constant number of adjacencies; the proof of the proposition shows $c_\Pi := 2$ suffices. Since all graphs with an odd cycle have at least one edge, the second condition is satisfied as well. For the last condition, consider a vertex-minimal graph G with an odd cycle; such a graph is Hamiltonian and has an odd number of vertices. Fix a minimum vertex cover X of G: walking along a Hamiltonian cycle of G, we must visit X before and after visiting every vertex from $V(G) \setminus X$ (as vertices of $V(G) \setminus X$ only have neighbors in X). Hence $|V(G) \setminus X| \leq X$ as every vertex in X is visited exactly one, and therefore $|V(G)| \leq 2|X|$ which proves that $p(n) := 2n$ suffices for the polynomial in Property (iii). We obtain a kernel with $\mathcal{O}(|X|^3)$ vertices.

CHORDAL DELETION (VC). A graph G is chordal if all its cycles of length at least four have a chord; this can be stated equivalently as saying that it does not contain a graph with a chordless cycle as an induced subgraph. If we take Π to be the class of graphs which have a chordless cycle, we can express CHORDAL DELETION (VC) as an instantiation of DELETION DISTANCE TO Π-FREE (VC). The proof of Proposition 2 shows the property is characterized by three adjacencies. As all graphs with a chordless cycle contain an edge, the second property is satisfied. Similarly as before, a vertex-minimal graph with a chordless cycle is Hamiltonian and hence $p(n) := 2n$ suffices for Property (iii). The resulting kernel has $\mathcal{O}(|X|^4)$ vertices.

Planarization (vc). Since this problem is a special case of $\mathcal{F}$-Minor-Free Deletion (vc) for $\mathcal{F} := \{K_5, K_{3,3}\}$, and both forbidden minors are nonempty, the proof given in the main text shows that this problem has a kernel with $\mathcal{O}(|X|^5)$ vertices.

η-Transversal (vc). Recall that the η-Transversal problem asks for a vertex set whose removal results in a graph of treewidth at most η (see Section E). Since treewidth does not increase when taking a minor, the class of graphs of treewidth at most η is closed under minors. By the famous results of Robertson and Seymour, this implies that for each η there is a finite obstruction set $\mathcal{F}_\eta$ such that G has treewidth at most η if and only if G avoids all graphs in $\mathcal{F}_\eta$ as a minor. If $\eta = 0$ then η-Transversal (vc) is equivalent to Vertex Cover (vc) which we handled earlier. It is easy to see that for $\eta > 0$ the minimal obstruction sets $\mathcal{F}_\eta$ do not contain empty graphs, as empty graphs have treewidth zero and cannot be obstructions to having treewidth $\eta > 0$. Hence we may obtain a polynomial kernel for η-Transversal (vc) by using the obstruction set $\mathcal{F}_\eta$ in the more general $\mathcal{F}$-Minor-Free Deletion (vc) scheme. The kernel size is $\mathcal{O}(|X|^{\Delta+1})$ where $\Delta := \max_{H \in \mathcal{F}_\eta} \Delta(H)$. □

C.4 Proof of Theorem 2

Proof (of Theorem 2). The kernelization reduces an instance (G, X, k) by executing Reduce $(G, X, p(|X|), c_\Pi)$ to obtain a graph G', and outputs the instance (G', X, k). By Observation 1 this can be done in polynomial time and results in a graph whose size is appropriately bounded; it remains to prove that the two instances are equivalent.

Since G' is an induced subgraph of G, any solution contained in G' is also contained in G: so if (G', X, k) is a yes-instance, then (G, X, k) is as well. Assume then that (G, X, k) is a yes-instance and let $P \subseteq V(G)$ be such that $G[P] \in \Pi$ and $|P| \geq k$. Clearly, $X \cap P$ is a vertex cover of $G[P]$, so $|P| \leq p(|X \cap P|) \leq p(|X|)$ by Property (ii). Since the reduction procedure is executed with a value $\ell := p(|X|)$ and $|P| \leq p(|X|)$, by applying Lemma 1 with an empty set for S we find that G' contains a set P' of the same size as P such that $G'[P] \in \Pi$. This proves that (G', X, k) is a yes-instance and shows the correctness of the kernelization. □

C.5 Corollary 2

Corollary 2. *All problems in Table 2 fit into the framework of Theorem 2 and admit polynomial kernels parameterized by the size of a given vertex cover.*

Proof (of Corollary 2). We consider the problems in the order of Table 2 and show how they fit into the framework.

Long Cycle (vc). Observe that if G has a cycle on k vertices $(v_1, \ldots, v_k)$ then the graph $G[\{v_1, \ldots, v_k\}]$ is Hamiltonian. So G has a k-cycle if and only if G has an induced Hamiltonian subgraph on k vertices. Hence Long Cycle (vc) is equivalent to Largest Induced Π-subgraph (vc) by letting Π be the class of Hamiltonian graphs. By Proposition 2 this class is characterized by two adjacencies. For all Hamiltonian graphs G' it holds that $|V(G')| \leq 2|\operatorname{VC}(G')|$, since for any Hamiltonian cycle C and vertex cover X, any vertex outside X must have its unique predecessor on the cycle in X. Hence Property (ii) is satisfied as well and we obtain a kernel with $\mathcal{O}(|X|^3)$ vertices. The proof for Long Path (vc) is analogous.

H-Packing (vc). A graph G admits an H-packing of k disjoint subgraphs, if and only if G has an induced subgraph on $k|V(H)|$ vertices which admits a perfect H-packing. Choosing Π as the graphs with a perfect H-packing therefore allows us to model the packing problem as an instantiation of Largest Induced Π-subgraph (vc), by scaling the target value k by a factor $|V(H)|$. Proposition 2 shows that Π is characterized by $\Delta(H)$ adjacencies. Let us now prove that the second condition is satisfied for this Π, by utilizing the fact that we demand H to be nonempty. Consider a graph G with a perfect H-packing for a nonempty H, and let X be a minimum vertex cover of G. Each subgraph in the packing contains at least one edge, so each subgraph in the packing has size $|V(H)|$ and contains a vertex from X. Hence $|V(G)| \leq |X| \cdot |V(H)|$, which proves that $p(n) := n \cdot |V(H)|$ suffices for the polynomial. For fixed H this results in a kernel with $\mathcal{O}(|X|^{\Delta(|H|)+1})$ vertices. ☐

D Omitted proofs for Section 4

D.1 Proofs for Testing for Cliques

Lemma 3. *Rule 3 is safe.*

Proof. As G' is a subgraph of G, any clique minor in G' also exists in G. Therefore, we need to argue that if G admits a K_t minor, then G' does as well.

Let ϕ be a clique minor model in G. If s does not belong to any branch set $\phi(v)$ for $v \in K_t$, then ϕ is also a clique minor in G' and we are done. Assume then that $s \in \phi(v)$. Observe that $\phi(v)$ has to contain at least one vertex from $N_G(s)$, as otherwise we would have that $\phi(v) = \{s\}$ and this $\phi(v)$ would be able to touch at most $t-2$ other branch sets. Obtain ϕ' from ϕ by removing s from $\phi(v)$ and observe that ϕ' is a K_t model in G': all the connections that were introduced by s are already present in the clique $N_G(s)$. ☐

Lemma 4. *If Reduction Rules 1-3 are not applicable, then* $|V(G)| \leq (|X|+1)^4$.

Proof. After exhausting Reduction Rules 2 and 3, there are no simplicial vertices in $V(G) \setminus X$. As Reduction Rule 1 is not applicable, for each of the at most $\binom{|X|}{2}$ non-edges in X there are at most $(|X|+1)^2$ vertices of $V(G) \setminus X$ adjacent to both endpoints. As every vertex of $V(G) \setminus X$ is adjacent to the endpoints of some non-edge, $|V(G')| \leq |X| + \binom{|X|}{2} \cdot (|X|+1)^2 \leq (|X|+1)^4$. ☐

D.2 Proofs of Testing for Bicliques

Polynomial Kernel for Induced $K_{c,t}$-testing We give a polynomial kernel for the following problem.

> Induced $K_{c,t}$ Subgraph Test (vc)
> **Input:** A graph G with vertex cover $X \subseteq V(G)$, and an integer $t \geq 1$.
> **Parameter:** The size $|X|$ of the vertex cover.
> **Question:** Does G contain $K_{c,t}$ as an induced subgraph?

Observe that c is treated as a constant, rather than a variable.

Theorem 5. Induced $K_{c,t}$ Subgraph Test (vc) *admits a polynomial kernel for every constant c.*

Proof. We may assume that $t > c$, as otherwise $K_{c,t}$ is a graph of constant size and we can solve the problem in polynomial time via brute-force.

Let (G, X, t) be the input instance. Whenever we assume that (G, X, t) is a yes-instance, we fix some induced $K_{c,t}$ subgraph of G and denote its bipartition by (A, B), where $|A| = c$ and $|B| = t$.

We provide a polynomial time algorithm that returns either:

(i) one instance of Induced $K_{c,t}$ Subgraph Test (vc) with $\mathcal{O}(|X|^{c+1})$ vertices that is equivalent to (G, X, t), or
(ii) at most $\binom{|X|}{c}$ instances of Independent Set, each with $\mathcal{O}(|X|^2)$ vertices, such (G, X, t) is a yes-instance if and only if at least one of them is a yes-instance.

In Case (i) we can simply output the obtained instance of Induced $K_{c,t}$ Subgraph Test (vc). In Case (ii) note that the total bitsize of returned instances is polynomial in $|X|$, and a nondeterministic Turing machine can verify in polynomial time whether at least one of them is a yes-instance. Therefore, as Induced $K_{c,t}$ Subgraph Test (vc) is NP-complete, we can use the reduction to create one instance of Induced $K_{c,t}$ Subgraph Test (vc) of size polynomial in $|X|$, which is a yes-instance if and only if at least one of the returned instances was. Thus we construct the kernel.

First, we exhaustively apply the following reduction rule. For every vertex $v \in V(G)$ we check, whether its neighborhood contains an independent set of size c. If this is not the case, we may safely delete this vertex as it cannot be contained in any induced $K_{c,t}$; note that in this step we use the assumption that $t > c$ to verify that the vertex cannot be in part A of the solution, either. This check can be done in polynomial time by iterating through all the subsets of $N_G(v)$ of size c. From now on we may assume that each vertex of the graph has an independent set of size c in its neighborhood.

Observe that if $|V(G) \setminus X| \geq t \cdot \binom{|X|}{c}$, then (G, X, t) is a yes-instance, as some t vertices of $V(G) \setminus X$ are adjacent to the same independent set of size c in X. In this case we output a trivial yes-instance. Moreover, if this is not the case but $t \leq |X|$, then $|V(G)| \leq |X| + |X| \cdot \binom{|X|}{c} = \mathcal{O}(|X|^{c+1})$ and we may output the graph obtained so far as the kernel in Case (i).

We are left with the case that $t > |X|$. Note that if (G, X, t) is a YES-instance, then part B has to contain at least one vertex from $V(G) \setminus X$, which means that $A \subseteq X$. For each subset $A' \subseteq X$ of size c that induces an independent set, we construct an instance $(G_{A'}, X_{A'}, t)$ of INDEPENDENT SET (VC), by taking $G_{A'} = G[\bigcap_{v \in A'} N_G(v)]$ and $X_{A'} = X \cap V(G_{A'})$. Observe that if (G, X, t) has a solution with $A = A'$, then $(G_{A'}, X_{A'}, t)$ is a YES-instance as the corresponding part B is contained in $\bigcap_{v \in A'} N_G(v)$. On the other hand, if $G_{A'}$ contains an independent set B' of size t, then $A' \cup B'$ induces a $K_{c,t}$ in G. Therefore, (G, X, t) is a YES-instance if and only if then at least one of the instances $(G_{A'}, X_{A'}, t)$ is a YES-instance. Observing that INDEPENDENT SET (VC) is equivalent to VERTEX COVER (VC) (by going to the dual target value $k' := n - k$, while keeping the parameter $|X|$ the same) we can use the kernelization algorithm for VERTEX COVER (VC) from Theorem 1 to every instance $(G_{A'}, X_{A'}, t)$, thus obtaining a sequence of instances of INDEPENDENT SET with $\mathcal{O}(|X|^2)$ vertices each, that can be returned in Case (ii). □

Kernel Lower Bound for Induced $K_{s,t}$-testing In this section we prove a kernel lower bound for the following problem.

INDUCED $K_{s,t}$ SUBGRAPH TEST (VC)
Input: A graph G with vertex cover $X \subseteq V(G)$, and integers $s, t \geq 1$.
Parameter: The size $|X|$ of the vertex cover.
Question: Does G contain $K_{s,t}$ as an induced subgraph?

Observe that the crucial difference with INDUCED $K_{c,t}$ SUBGRAPH TEST (VC) is that the value s is part of the input, rather than a constant. We base our cross-composition on the balanced biclique problem in bipartite graphs.

BALANCED BICLIQUE IN BIPARTITE GRAPHS
Input: A bipartite graph G with partite sets $A \cup B$, and an integer $k \geq 1$.
Question: Are there subsets $S \subseteq A$ and $T \subseteq B$ such that $G[S \cup T]$ is a biclique, and $|S| = |T| = k$?

Theorem 6. INDUCED $K_{s,t}$ SUBGRAPH TEST (VC) *does not admit a polynomial kernel unless NP $\subseteq$ coNP/poly.*

Proof. We prove that BALANCED BICLIQUE IN BIPARTITE GRAPHS cross-composes into INDUCED $K_{s,t}$ SUBGRAPH TEST (VC), which suffices to establish the claim by Theorem 4 and the NP-completeness of the classical problem (GT24 in Garey and Johnson). Define a polynomial equivalence relation $\mathcal{R}$ as follows. Two strings in Σ^* are equivalent if (a) they both encode malformed instances, or (b) they encode valid instances (G_1, A_1, B_1, k_1) and (G_2, A_2, B_2, k_2) of BALANCED BICLIQUE IN BIPARTITE GRAPHS such that $|A_1| = |A_2|$, $|B_1| = |B_2|$ and $k_1 = k_2$. This relationship $\mathcal{R}$ partitions a set of instances on at most n vertices each into $\mathcal{O}(n^3)$ equivalence classes, and is therefore a polynomial equivalence relationship.

We compose instances which are equivalent under $\mathcal{R}$. So the input consists of r instances $(G_0, A_0, B_0, k), \ldots, (G_{r-1}, A_{r-1}, B_{r-1}, k)$ of Balanced Biclique in Bipartite Graphs which all agree on the number of vertices in each partite set, and on the value of k. By duplicating some instances we may assume without loss of generality that r is a power of two. Let $n := |B_1| = \ldots = |B_r|$ and $m := |A_1| = \ldots = |A_r|$. For $i \in [r]$ label the vertices in B_i as $b_{i,1}, \ldots, b_{i,n}$. We build a graph G^* with vertex cover X^* as follows.

- Initialize G^* as the disjoint union of the input graphs $G_1, \ldots, G_r$.
- For each $j \in [n]$, identify the vertices $b_{1,j}, \ldots, b_{r,j}$ into a single vertex b^*_j. Let $B^* := \{b^*_1, \ldots, b^*_n\}$ contain the resulting vertices, and observe that at this stage in the construction $G^*[A_i \cup B^*]$ is isomorphic to G_i for $i \in [r]$.
- For $j \in [\log r]$, add to G^* a biclique C_j isomorphic to $K_{n+1,n+1}$, with partite sets denoted by P_j and Q_j, $|P_j| = |Q_j| = n + 1$. The set of vertices corresponding to one value of j will be called the *bit selector* of j as it will be used in valid solutions to select the bitvalue of the binary representation of the input instance corresponding to this solution.
- For $j \in [\log r]$, make the vertices of P_j adjacent to the vertices of A_i if the j-th bit in the binary representation of number i is a one. Similarly, make the vertices Q_j adjacent to A_i if the j-th bit of i is a zero.
- Add a set D of $(n+1)(1+2\log r)$ vertices, adjacent to all the vertices of B^* and all the vertices of all the bit selectors.
- Let X^* contain the vertices of B^* and all the vertices of all the bit selectors. Observe that $|X^*| = n + 2(n+1)\log r = |D| - 1$ and that $G^* - X^*$ is an independent set containing all the sets A_i and the set D; hence X^* is a vertex cover of G^* whose size is suitably bounded for a cross-composition.

The construction is completed by setting $s = k + (n+1)\log r$ and $t = k + (n+1)(1+2\log r) = k + |D|$. We now prove the completeness and soundness of the composition via two claims.

Claim. If for some i the instance (G_i, A_i, B_i, k) is a yes-instance of Balanced Biclique in Bipartite Graphs, then (G^*, X^*, s, t) is a yes-instance of Induced $K_{s,t}$ Subgraph Test (vc).

Proof. Let $S \subseteq A_i$ and $T \subseteq B_i$ be such that $|S| = |T| = k$ and $G_i[S \cup T]$ is a biclique. Let T' be the image of T in the identifications, i.e., $T' = \{b^*_j \mid b_{i,j} \in T\}$. For $j \in [\log r]$ define $R_j := P_j$ if the j-th bit of binary encoding of i is equal to one, and define $R_j := Q_j$ otherwise. We claim that the set $S \cup T' \cup D \cup \bigcup_{j \in [\log r]} R_j$ induces a biclique in G^*, with $T' \cup \bigcup_{j \in [\log r]} R_j$ as one partite set and $S \cup D$ as the second. Indeed, observe that

- $D, S, T', \bigcup_{j \in [\log r]} R_j$ are independent sets by the construction of G^*;
- there is no edge between D and S;
- there is no edge between $\bigcup_{j \in [\log r]} R_j$ and T';

– D is adjacent to the whole set X^*, so in particular to $T' \cup \bigcup_{j \in [\log r]} R_j$;
– as $S \subseteq A_i$, by the construction of G^* we have that every vertex of S is adjacent to every vertex of R_j, for all $j \in [\log r]$;
– all vertices in S are adjacent to all vertices of T', as $G_i[S \cup T]$ is a biclique and $G^*[A_i \cup B^*]$ is isomorphic to G_i.

We conclude the proof by checking that $|T' \cup \bigcup_{j \in [\log r]} R_j| = k + (n+1) \log r = s$ and $|S \cup D| = k + |D| = t$. □

Claim. If (G^*, X^*, s, t) is a YES-instance of INDUCED $K_{s,t}$ SUBGRAPH TEST (VC), then for some i the instance (G_i, A_i, B_i, k) is a YES-instance of BALANCED BICLIQUE IN BIPARTITE GRAPHS.

Proof. Assume that there exist sets S^* and T^*, $|S^*| = s$ and $|T^*| = t$, such that $G^*[S^* \cup T^*]$ is a biclique with S^* and T^* as partite sets. As $|T^*| = t \geq |D| > |X^*|$, the set $T^* \setminus X^*$ is nonempty. This means that in T^* there is a vertex with the whole neighborhood entirely contained in X^*, so $S^* \subseteq X^*$. From every pair (P_j, Q_j), for $j \in [\log r]$, the independent set S^* can have a nonempty intersection with at most one of them. Assume that for some j we have $P_j \cap S^* = Q_j \cap S^* = \emptyset$. It follows that $|S^*| \leq (n+1) \log r - (n+1) + |B^*| < (n+1) \log r + k = s$, which is a contradiction. Hence, for all $j \in [\log r]$ set S^* has a nonempty intersection with exactly one set of P_j and Q_j. Moreover, observe that $|S^* \cap \bigcup_{j \in [\log r]} (P_j \cup Q_j)| \leq (n+1) \log r$, so $|S^* \cap B^*| \geq k$.

Define i as an integer with $\log r$ binary digits, such that the j-th bit is equal to one if $P_j \cap S^* \neq \emptyset$ and is equal to zero if $Q_j \cap S^* \neq \emptyset$. By the construction of G^*, the set A_i is the only set from $\{A_0, \ldots, A_{r-1}\}$ which contains vertices simultaneously adjacent to all vertices from S^* contained in bit selectors. There is no edge between B^* and bit selectors, so we infer that $T^* \subseteq A_i \cup D$. As $|T^*| = k + |D|$, we infer that $|T^* \cap A_i| \geq k$.

Recall that $G^*[B^* \cup A_i]$ is isomorphic to G_i, hence $(T^* \cap A_i) \cup (S^* \cap B^*)$ induces a biclique in a graph isomorphic to G_i. As $|T^* \cap A_i|, |S^* \cap B^*| \geq k$, we infer that G_i is a YES-instance of BALANCED BICLIQUE IN BIPARTITE GRAPHS. □

As this proves that the output instance acts as the logical OR of the inputs, it concludes the cross-composition and proves a kernel lower bound by Theorem 4. □

D.3 Proofs for Testing for Paths

We prove in this section that the following problem is unlikely to admit polynomial kernels.

LONG INDUCED PATH (VC)
Input: A graph G with vertex cover $X \subseteq V(G)$, and an integer $k \geq 1$.
Parameter: The size $|X|$ of the vertex cover.
Question: Is there a set $S \subseteq V(G)$ of size at least k such that $G[S]$ is a simple path?

Using cross-composition, we start from the following classical problem.

> HAMILTONIAN $s-t$ PATH
> **Input:** A graph G with distinct vertices s and t.
> **Question:** Is there a Hamiltonian path from s to t in G?

The following simple proposition will be of use during the proof.

Proposition 3. *If G is a graph with vertex cover $X \subseteq V(G)$ then any simple path in G contains at most $2|X|+1$ vertices.*

Proof. Let P be a path in G, and let $y_1, y_2, \ldots, y_p$ be the vertices from $G-X$ on this path in the order in which they are visited. Since vertices of $G-X$ only have neighbors in X, path P must visit at least one vertex from X between visiting y_i and y_{i+1}, for $i \in [p-1]$. Because each vertex can be visited at most once it follows that P contains at most $|X|+1$ vertices from $G-X$, and since there are $|X|$ vertices in X the bound follows. □

Armed with these ingredients we can state and prove the kernel lower bound.

Theorem 7. LONG INDUCED PATH (VC) *does not admit a polynomial kernel unless NP $\subseteq$ coNP/poly.*

Proof. By Theorem 4 and the NP-completeness of HAMILTONIAN $s-t$ PATH (GT39 in Garey and Johnson), it is sufficient to show that HAMILTONIAN $s-t$ PATH cross-composes into LONG INDUCED PATH (VC). We define a polynomial equivalence relation $\mathcal{R}$ as follows. We say that two strings in Σ^* are equivalent if (a) they both encode malformed instances, or (b) they encode valid instances (G_1, s_1, t_1) and (G_2, s_2, t_2) of HAMILTONIAN $s-t$ PATH such that $|V(G_1)| = |V(G_2)|$. This implies that $\mathcal{R}$ partitions a set of instances on at most n vertices each into $\mathcal{O}(n)$ equivalence classes, and is therefore a polynomial equivalence relationship.

We show how to compose a set of instances which are equivalent under $\mathcal{R}$. So the input consists of r instances $(G_1, s_1, t_1), \ldots, (G_r, s_r, t_r)$ of HAMILTONIAN $s-t$ PATH such that $|V(G_i)| = n$ for $i \in [r]$. We may assume that $n \geq 9$, since we can solve smaller instances in constant time, reducing to a constant-size YES- or NO-instance. For $i \in [r]$ label the vertices in $V(G_i)$ as $v_1, \ldots, v_n$ such that $s_i = v_1$ and $t_i = v_n$. We build a graph G^* with vertex cover X^* as follows.

1. Add three simple paths P_A, P_B and P_C to G^*, containing n^3 vertices each. Let the endpoints of these paths be x_A, y_A, x_B, y_B and x_C, y_C respectively.
2. For $j \in [n]$ add a vertex v^*_j to G^*.
3. For $\{j,h\} \in \binom{[n]}{2}$ add a vertex $e_{j,h}$ to G^* and make it adjacent to v^*_j and v^*_h.
4. For $i \in [r]$, do the following. Add a vertex z_i to G^*. For all pairs $\{j,h\} \in \binom{[n]}{2}$ such that $\{v_j, v_h\} \notin E(G_i)$ add the edge $\{z_i, e_{j,h}\}$ to G^*.
5. Make y_A and y_B adjacent to all vertices z_i for $i \in [r]$.
6. Make x_B adjacent to v^*_1, and make x_C adjacent to v^*_n. This concludes the construction of G^*, which is illustrated in Fig. 1.

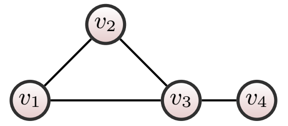

(a) Input instance (G_1, v_1, v_4) of Hamiltonian $s - t$ Path.

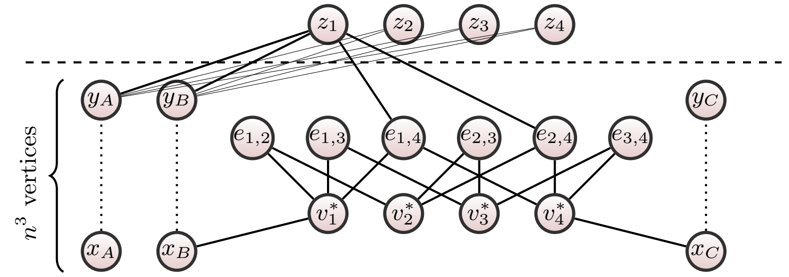

(b) Output instance of Long Induced Path (vc).

Fig. 1. An example of the lower-bound construction of Theorem 7. (a) The first input instance. (b) The graph G^* which is the result of cross-composing four inputs on $n = 4$ vertices each. Note that the composition algorithm only builds the graph G^* when $n \geq 9$, but this picture gives the correct intuition for the construction. Edges between $\{z_2, z_3, z_4\}$ and $\{e_{j,h} \mid \{j,h\} \in \binom{[4]}{2}\}$ have been omitted for readability. The vertices below the horizontal dashed line form the vertex cover X^*.

We define a set $X^* := V(G^*) \setminus \{z_i \mid i \in [r]\}$. Since we did not add any edges between the z-vertices, they form an independent set and therefore X^* is a vertex cover of G^*. It is easy to verify that the size of X^* is polynomial in n, and therefore the size of the parameter $|X^*|$ is suitably bounded for a cross-composition. We set $k^* := 3n^3 + 2n$. The construction can be performed in polynomial time, so it remains to prove that (G^*, X^*, k^*) is yes if and only if one of the input instances is yes. We first establish some properties of the constructed instance.

Claim. Let $S^* \subseteq V(G)$ induce a simple path in G^*, and let $P^* := G^*[S^*]$.

1. $|S^* \setminus (V(P_A) \cup V(P_B) \cup V(P_C))| < n^3 - 3$.
2. If there is a path $P_W \in \{P_A, P_B, P_C\}$ such that $|S^* \cap V(P_W)| \leq 3$ then $|S^*| < k^*$.

Proof. Define $\hat{G^*} := G^* - (V(P_A) \cup V(P_B) \cup V(P_C))$.

(1) For each of the paths P_A, P_B, P_C there are at most two vertices on the path which have neighbors outside the path. Hence if the take the path P^*, then deleting the vertices of $V(P_A)$ from P^* splits the path into at most three pieces, increasing the number of connected components by at most two. This also holds for P_B and P_C. Hence $\hat{P^*} := P^* - (V(P_A) \cup V(P_B) \cup V(P_C))$ is an

induced linear forest in $\hat{G}^*$ containing no more than seven connected components, with $|S^*| = |V(\hat{P}^*)|$. Each connected component of $\hat{P}^*$ is an induced path in $\hat{G}^*$. Since the set $\hat{X}^* := \{v_i^* \mid i \in [n]\} \cup \{e_{j,h} \mid \{j,h\} \in \binom{[n]}{2}\}$ is a vertex cover for $\hat{G}^*$ of size $n + \binom{n}{2}$ it follows by Proposition 3 that each connected component of $\hat{P}^*$ has at most $2n + 2\binom{n}{2} + 1$ vertices. Since the number of connected components is at most seven, the number of vertices in $\hat{P}^*$ is at most $7 \cdot (2n + 2\binom{n}{2} + 1)$, which is less than $n^3 - 3$ for $n \geq 9$.

(2) Assume that S^* contains at most three vertices from P_A; the other two cases will be completely analogous.

$$\begin{aligned}
|S^*| =& |S^* \setminus (V(P_A) \cup V(P_B) \cup V(P_C))| + \\
& |S^* \cap V(P_A)| + |S^* \cap V(P_B)| + |S^* \cap V(P_C)| \\
<& (n^3 - 3) + |S^* \cap V(P_A)| + \\
& |S^* \cap V(P_B)| + |S^* \cap V(P_C)| && \text{By (1).} \\
\leq& (n^3 - 3) + 3 + |S^* \cap V(P_B)| + |S^* \cap V(P_C)| && \text{By assumption.} \\
\leq& (n^3 - 3) + 3 + n^3 + n^3 && \text{By definition of } G^*. \\
\leq& k^*. && \text{By definition of } k^*.
\end{aligned}$$

Hence if there is one path among $\{P_A, P_B, P_C\}$ such that S^* contains at most three vertices on it, then $|S^*| < k^*$. □

We now prove that (G^*, X^*, k^*) indeed acts as the OR of the input instances. For the first direction, assume that G^* has a path on at least k^* vertices induced by the vertex set S^*. Let $P^* := G^*[S^*]$ be the path induced by S^*. By (2) the set S^* contains at least three vertices on each of the paths P_A, P_B, P_C. Since P_A and P_C each contain exactly one vertex which has neighbors outside the path, it is easy to see that $S^* \cup V(P_A) \cup V(P_C)$ is also an induced path; hence we may assume without loss of generality that S^* contains all vertices of P_A and P_C, which means that the endpoints of P^* must be the vertices x_A and y_C since they have degree one in G^*. Since no endpoint of P^* can lie on P_B, and S^* contains at least three vertices on P_B, it follows that S^* must contain all vertices of P_B since the internal vertices on that path do not have neighbors outside the path. Hence $V(P_A) \cup V(P_B) \cup V(P_C) \subseteq S^*$. Since the only neighbors of vertex y_A are the vertices z_i for $i \in [r]$ and the single neighbor on the path P_A, the path P^* must contain an edge $\{y_A, z_{i^*}\}$ for some $i^* \in [r]$ since y_A must have two neighbors on the path. By construction of G^* we know that $\{z_{i^*}, y_B\}$ is an edge in G^*. This implies that if we traverse the path P^* starting from the endpoint x_A then we traverse P_A, visit z_{i^*}, and then go to y_B. Since all vertices of P_B are in S^* the path then traverses P_B until it reaches x_B. The unique neighbor v_1^* of x_B not on P_B must be the successor of x_B on the path P^*. The path now visits some more vertices. Since all vertices of P_C are contained in S^*, and x_C is the only vertex of P_C adjacent to vertices not on P_C, the path P^* must finish by reaching x_C and traversing P_C.

Let us now consider the subpath P_{i^*} of P^* which starts at the successor of x_B on the path, and stops with the predecessor of x_C on the path. The successor of x_B must be v_1^*, and the predecessor of x_C must be v_n^*, since vertices x_B and x_C have degree two in G^*; hence P_{i^*} is an induced path from v_1^* to v_n^*. Since there are $2n^3+1$ vertices on P^* before x_B, and n^3 vertices on P^* on the final part from x_C to the endpoint, the subpath P_{i^*} must contain at least $k^* - 3n^3 - 1 = 2n - 1$ vertices. Since the vertices y_A and y_B are contained in S^* and are adjacent to z_{i^*}, the set S^* cannot contain any other vertices adjacent to z_{i^*} (otherwise these would induce an edge not on the path P^*). This implies that in particular, S^* cannot contain vertices $e_{j,h}$ for which $\{v_j, v_h\} \notin E(G_{i^*})$ since these were made adjacent to z_{i^*} in the construction. The set S^* cannot contain any vertices z_i for $i \neq i^*$, since all such vertices are adjacent to $y_A, y_B \in S^*$ and together with $z_{i^*} \in S^*$ such a vertex z_i would induce a cycle. This shows that the subpath P_{i^*} can contain only vertices v_j^* for $j \in [n]$, and vertices $e_{j,h}$ for $\{v_j, v_h\} \in E(G_{i^*})$. Since the edge vertices $e_{j,h}$ are only adjacent to the vertices which form their endpoints, it now follows that the edge set $\{\{v_j, v_h\} \mid e_{j,h} \in S^*\}$ is a path in G_{i^*} between $v_1 = s_{i^*}$ and $v_n = t_{i^*}$ containing $n-1$ edges and n vertices, which implies that G_{i^*} has a Hamiltonian $v_1 - v_n$ path and proves that G_{i^*} is YES.

For the reverse direction, assume that the set $C_{i^*} \subseteq E(G_{i^*})$ are the edges on a Hamiltonian $v_1 - v_n$ path in G_{i^*}. Then it is straightforward to verify using the construction of G^* that $S^* := V(P_A) \cup V(P_B) \cup V(P_C) \cup \{z_{i^*}\} \cup \{v_i^* \mid i \in [n]\} \cup \{e_{j,h} \mid \{v_j, v_h\} \in C_{i^*}\}$ induces a simple path in G^* and has size at least k^*. This concludes the proof. □

D.4 Proofs for Testing for Matchings

Recall that an *induced matching* in a graph G is a matching $Y \subseteq E(G)$ such that no edge in $E(G) \setminus Y$ connects the endpoints of two edges of Y, or equivalently, such that all connected components of the subgraph induced by the endpoints of Y are isomorphic to K_2. The *size* of an induced matching is measured in terms of the number of edges in it. The goal of this section is to prove a superpolynomial kernel lower bound for the following problem.

> Maximum Induced Matching (vc)
> **Input:** A graph G with vertex cover $X \subseteq V(G)$, and an integer $k \geq 1$.
> **Parameter:** The size $|X|$ of the vertex cover.
> **Question:** Is there an induced matching $Y \subseteq E(G)$ in G of size at least k?

Using the technique of cross-composition, we start from the following related classical problem.

> Maximum Induced Matching in Bipartite Graphs
> **Input:** A bipartite graph G with partite sets $A \cup B$, and an integer $k \geq 1$.
> **Question:** Is there an induced matching $Y \subseteq E(G)$ in G of size at least k?

The cross-composition uses a bit masking scheme, together with a repeated structure in the graph to simulate heavy-weight edges, to embed the OR of bipartite instances into a single bipartite instance of the parameterized problem with a small parameter value.

Theorem 8. MAXIMUM INDUCED MATCHING (VC) *does not admit a polynomial kernel unless NP $\subseteq$ coNP/poly.*

Proof. We prove that MAXIMUM INDUCED MATCHING IN BIPARTITE GRAPHS cross-composes into MAXIMUM INDUCED MATCHING (VC), which suffices to establish the claim by Theorem 4 and the NP-completeness of the classical problem [8]. Define a polynomial equivalence relation $\mathcal{R}$ as follows. Two strings in Σ^* are equivalent if (a) they both encode malformed instances, or (b) they encode valid instances (G_1, A_1, B_1, k_1) and (G_2, A_2, B_2, k_2) of MAXIMUM INDUCED MATCHING IN BIPARTITE GRAPHS such that $|A_1| = |A_2|$, $|B_1| = |B_2|$ and $k_1 = k_2$. This relationship $\mathcal{R}$ partitions a set of instances on at most n vertices each into $\mathcal{O}(n^3)$ equivalence classes, and is therefore a polynomial equivalence relationship.

We compose instances which are equivalent under $\mathcal{R}$. So the input consists of r instances $(G_1, A_1, B_1, k), \ldots, (G_r, A_r, B_r, k)$ of MAXIMUM INDUCED MATCHING IN BIPARTITE GRAPHS which all agree on the number of vertices in each partite set, and on the value of k. By duplicating some instances we may assume without loss of generality that r is a power of two. Let $n := |B_1| = \ldots = |B_r|$. For $i \in [r]$ label the vertices in B_i as $b_{i,1}, \ldots, b_{i,n}$. We build a graph G^* with vertex cover X^* as follows.

- Initialize G^* as the disjoint union of the input graphs $G_1, \ldots, G_r$.
- For each $j \in [n]$, identify the vertices $b_{1,j}, \ldots, b_{r,j}$ into a single vertex b^*_j. Let $B^* := \{b^*_1, \ldots, b^*_n\}$ contain the resulting vertices, and observe that at this stage in the construction $G^*[A_i \cup B^*]$ is isomorphic to G_i for $i \in [r]$.
- For $j \in [\log r]$, add vertices $\{x_{s,j}, y_{s,j}, z_{s,j} \mid s \in [n]\}$ to G^* and turn each triplet into a clique. The set of vertices corresponding to one value of j will be called the *bit selector* of j as it will be used in valid solutions to select the bitvalue of the binary representation of the input instance corresponding to this solution.
- For $j \in [\log r]$, make the vertices $\{x_{s,j} \mid s \in [n]\}$ adjacent to the vertices A_i if the j-th bit in the binary representation of number i is a one. Similarly, make the vertices $\{y_{s,j} \mid s \in [n]\}$ adjacent to A_i if the j-th bit of i is a zero. Let $\{\{x_{s,j}, z_{s,j}\} \mid s \in [n]\}$ be the *x-edges of position j*, and let $\{\{y_{s,j}, z_{s,j}\} \mid s \in [n]\}$ be the *y-edges of position j*.
- Let X^* contain the vertices of B^* and all the vertices of all the bit selectors. Observe that $|X^*| = n + 3n \log r$ and that $G^* - X^*$ is an independent set containing all the sets A_i; hence X^* is a vertex cover of G^* whose size is suitably bounded for a cross-composition.

The construction is completed by setting $k^* := k + n \log r$ and using the instance (G^*, X^*, k^*) as the output of the cross-composition. We will need the following structural claim.

Claim. G^* has a maximum induced matching $Y^* \subseteq E(G^*)$ such that for every $j \in [\log r]$ and $s \in [n]$, if Y^* contains an edge incident on the triple $\{x_{s,j}, y_{s,j}, z_{s,j}\}$ then Y^* contains $\{x_{s,j}, z_{s,j}\}$ or $\{y_{s,j}, z_{s,j}\}$.

Proof. Suppose Y^* is a maximum induced matching containing an edge incident on the triple $\{x_{s,j}, y_{s,j}, z_{s,j}\}$ for some choice of j and s, but the edge is neither $\{x_{s,j}, z_{s,j}\}$ nor $\{y_{s,j}, z_{s,j}\}$. As the triple forms a clique in G^*, by the induced property it follows that Y^* contains exactly one edge incident to it. Since vertex $z_{s,j}$ is only adjacent to $x_{s,j}$ and $y_{s,j}$ we find that the edge $e \in Y^*$ incident on the triple, is incident on at least one of the vertices $x_{s,j}$ or $y_{s,j}$. If e is incident on $x_{s,j}$ then we may replace it by the edge $\{x_{s,j}, z_{s,j}\}$ to obtain another induced matching; $z_{s,j}$ was not matched before, and is not adjacent to any matched vertices except $x_{s,j}$. Similarly we may replace e by $\{y_{s,j}, z_{s,j}\}$ if e is incident on $y_{s,j}$. As this replacement step can be performed independently for each triple, the claim follows. □

Claim. G^* has a maximum induced matching $Y^* \subseteq E(G^*)$ such that for every bit position $j \in [\log r]$, either all the x-edges of position j are in Y^*, or all the y-edges of position j are in Y^*.

Proof. Consider a maximum induced matching Y^* in G^*, and assume there is some bit position $j \in [\log r]$ for which the claim does not hold. By the previous claim we may assume that if Y^* contains an edge incident on a triple $\{x_{s,j}, y_{s,j}, z_{s,j}\}$, then it is the x-edge $\{x_{s,j}, z_{s,j}\}$ or the y-edge $\{y_{s,j}, z_{s,j}\}$.

If at least one x-edge (resp. y-edge) of position j is contained in Y^*, then it is easy to verify that removing all edges incident on the vertices of bit selector j and adding all x-edges (resp. y-edges) for that bit selector results in an induced matching which is not smaller, and in which the status of edges for other bit selectors is not changed; this follows from the fact that the adjacencies of the respective vertices to the outside the bit selector are identical. So in the remainder it suffices to consider a bit position $j \in [\log r]$ for which Y^* contains no edge incident on a vertex in the bit selector. We will exhibit an induced matching which is at least as large as Y^* and which has the desired form.

Observe that the bit selector for position j contains n triples, each of which forms a clique. As an induced matching cannot contain two edges incident on the same clique, Y^* contains at most one edge incident on each triple for each bit selector $j' \neq j$, and by assumption it contains no edges incident on bit selector j. Since the union of the sets A_i for $i \in [r]$ forms an independent set in G^*, all matching edges in Y^* have at least one endpoint in B^*, or one endpoint in a bit selector. As B^* has exactly n vertices, this bounds the number of edges in Y^* by $n + ((\log r) - 1)n = n \log r$. Now observe that the union of all the x-edges of the bit selectors forms an induced matching of size $n \log r$, and has the desired form. As we assumed Y^* to be maximum, the described induced matching is also maximum which concludes the proof. □

To complete the cross-composition it remains to prove that the constructed instance acts as the OR of the inputs. For the first direction, assume that G^*

has a maximum induced matching $Y^* \subseteq E(G^*)$ of size at least k^*. By the second claim we may assume that for each $j \in [\log r]$, the matching Y^* contains all the x-edges or all the y-edges of position j. Now consider the instance number i^* whose binary expansion has a zero (resp. one) in the j-th bit position if Y^* contains the x-edges (resp. y-edges) of bit selector j. By definition of the adjacencies of the bit selectors it follows that for all instance numbers $i' \in [r]$ with $i' \neq i^*$, no vertex of $A_{i'}$ is the endpoint of an edge in Y^*. To see this, consider a bit position $j \in [\log r]$ where the binary expansion of i' and i^* differ; the x-vertices (resp. y-vertices) of instance selector j are endpoint of edges in Y^* whose other endpoints are formed by the z-vertices. As the x-vertices (resp. y-vertices) are adjacent to $A_{i'}$ by the choice of j, inducedness of the matching shows that $A_{i'}$ contains no endpoints of matching edges. Hence Y^* is also an induced matching, of the same size, in the graph obtained from G^* by removing the vertices $A_{i'}$ for $i' \neq i^*$. Each triple of an instance selector is a clique, and by assumption on the form of Y^* the matching contains the x-edge or the y-edge of the triple. Since an induced matching cannot contain two edges incident on the same clique, this shows that no edges between A_{i^*} and an instance selector can be contained in Y^*. Therefore it follows that if we delete the vertices $A_{i'}$ for $i' \neq i^*$ together with the vertices of the instance selectors from G^*, we are left with an induced submatching of size at least $k' - (n \log r) = k$. But the resulting graph is $G^*[B^* \cup A_{i^*}]$, and as observed in the construction of G^* it is isomorphic to G_{i^*}, which proves that G_{i^*} contains an induced matching of size k and is a YES-instance.

For the reverse direction, assume there is some index $i^* \in [r]$ such that G_{i^*} has an induced matching Y of size k. As $G^*[B^* \cup A_{i^*}]$ is isomorphic to G_{i^*}, this implies that the induced subgraph admits an induced matching of size k. Now augment this into an induced matching in G^* by adding the x-edges of the bit selectors for positions j where the binary expansion of i^* has a zero, and the y-edges where the expansion has a one. Using the description of G^* it is easy to verify that the resulting set of edges is an induced matching, containing a total of $k + n \log r$ edges. This proves that (G^*, X^*, k^*) is a YES-instance.

As the construction can be carried out in polynomial time and embeds the OR of the input instances into a single instance of the target problem with parameter value $|X^*| = n + 3n \log r$, this concludes the proof of Theorem 8. □

E Problem Compendium

For completeness we provide a definition for all problems mentioned in Tables 1 and 2.

VERTEX COVER (VC)
Input: A graph G with vertex cover $X \subseteq V(G)$, and an integer $k \geq 1$.
Parameter: The size $|X|$ of the vertex cover.
Question: Does G have a vertex cover of size at most k, i.e., is there a set $S \subseteq V(G)$ of size at most k such that $G - S$ is an empty graph?

The preceding problem may appear strange at first; one should think of it as preprocessing the VERTEX COVER problem when an approximate solution is supplied to the preprocessing algorithm. It is strongly related to the natural parameterization k-VERTEX COVER.

ODD CYCLE TRANSVERSAL (VC)
Input: A graph G with vertex cover $X \subseteq V(G)$, and an integer $k \geq 1$.
Parameter: The size $|X|$ of the vertex cover.
Question: Is there a set $S \subseteq V(G)$ of size at most k such that $G - S$ is bipartite?

PLANARIZATION (VC)
Input: A graph G with vertex cover $X \subseteq V(G)$, and an integer $k \geq 1$.
Parameter: The size $|X|$ of the vertex cover.
Question: Is there a set $S \subseteq V(G)$ of size at most k such that $G - S$ is planar?

η-TRANSVERSAL (VC)
Input: A graph G with vertex cover $X \subseteq V(G)$, and an integer $k \geq 1$.
Parameter: The size $|X|$ of the vertex cover.
Question: Is there a set $S \subseteq V(G)$ of size at most k such that $G - S$ has treewidth at most η?

$\mathcal{F}$-MINOR-FREE DELETION (VC)
Input: A graph G with vertex cover $X \subseteq V(G)$, and an integer $k \geq 1$.
Parameter: The size $|X|$ of the vertex cover.
Question: Is there a set $S \subseteq V(G)$ of size at most k such that $G - S$ does not contain any graph in $\mathcal{F}$ as a minor? ($\mathcal{F}$ can be any fixed finite set of nonempty graphs.)

Observe that if $\mathcal{F}$ contains an empty graph on ℓ vertices, then any graph with ℓ or more vertices contains an $\mathcal{F}$ minor. Hence in such cases the set of $\mathcal{F}$-minor-free graphs is finite which allows $\mathcal{F}$-MINOR-FREE DELETION to be solved in polynomial time: we can simply try all vertex subsets of less than ℓ vertices and test whether they induce one of the $\mathcal{F}$-minor-free graphs, resulting in a runtime of $\mathcal{O}(n^{\ell+\mathcal{O}(1)})$ which is polynomial for fixed $\mathcal{F}$.

LONG CYCLE (VC)
Input: A graph G with vertex cover $X \subseteq V(G)$, and an integer $k \geq 1$.
Parameter: The size $|X|$ of the vertex cover.
Question: Does G contain a simple cycle on at least k vertices?

LONG PATH (VC)
Input: A graph G with vertex cover $X \subseteq V(G)$, and an integer $k \geq 1$.
Parameter: The size $|X|$ of the vertex cover.
Question: Does G contain a simple path on at least k vertices?

H-Packing (vc)
Input: A graph G with vertex cover $X \subseteq V(G)$, and an integer $k \geq 1$.
Parameter: The size $|X|$ of the vertex cover.
Question: Does G contain at least k vertex-disjoint subgraphs isomorphic to H? (H can be any fixed nonempty graph.)

Observe that the well-known Triangle Packing problem is a special case of the previous problem.